10/586910
IAP11 Rec'd PCT/PTO 24 JUL 2006

AUTOMATIC ALLOCATION OF A NETWORK NUMBER FOR A COMMUNICATIONS ROUTER IN AN IPV6 NETWORK

The present invention relates to communications networks of the Internet type and more specifically to communications networks based on the IPv6 (*Internet Protocol - version 6*) protocol stack. It relates more particularly to automatically configuring such a network.

A network of the above type comprises a set of routers (also called "network equipments"), whose function is to route data traffic between a sender and a receiver. Each router has one or more interfaces and uses each interface to communicate with one or more other routers.

In accompanying Figure 1, the router R_A has two interfaces I_{A1} and I_{A2}. It uses the interface I_{A2} to communicate with a single router R_D via the interface I_D of that single router. It uses the interface I_{A1} to communicate with two routers R_B and R_C connected to the same link via their respective interfaces I_B and I_C. In reality, a router or network equipment has at least two or three interfaces. It must therefore "switch" each incoming data packet from one of its interfaces to one or the other of its other interfaces. This choice is by no means trivial. The choice mechanism is known as "routing".

In order to be able to route traffic correctly from one point to another, each router has access to a routing table giving the correspondence between a set of addresses and an output interface: a router receiving a data packet having a particular destination address is therefore able to determine to which of its interfaces it is to send it.

Before a communications network can route data traffic in this way, it is therefore necessary to configure the network, in particular by allocating addresses to each of the interfaces of each of the routers of the network, after which the routing tables

are constructed.

The specifications of an IPv6 network are set out in the document RFC 2460 published by the IETF (Internet Engineering Task Force). The document "IP Version 6 Addressing Architecture" (draft-ieff=ipv6-addr-v4=00.txt) specifies more particularly how the addresses of these interfaces are to be constituted.

There are several address types, but the present invention relates to "global unicast" addresses, which uniquely identify an interface of a communications router in a network. In the case of the Internet, which connects a plurality of sub-networks worldwide, uniqueness is required at the global level. For clarity, these "global unicast" addresses are referred to below simply as "global" addresses.

A global address primarily comprises a first part, typically comprising 64 bits, and a second part, typically comprising 64 bits, the global address therefore comprising 128 bits in total.

The second part is constituted from a unique identifier of the interface, in a manner which is specified in section 2.5.1 of the document "IP Version 6 Addressing Architecture". It may be constituted from an universal identifier, for example of the type defined in the IEEE 802 MAC (Media Access Control) standard or the IEEE EUI-64 (Extended Universal Identifier) standard.

Each router can easily determine this second part autonomously and automatically.

However, there is no automatic method for a router to determine the first part, which is usually called the "network number".

At present, in IPv6 networks, this part is determined manually by an operator responsible for configuring the network. The operator connects to each router in order to allocate it a global address for each interface, ideally in accordance with an optimized addressing scheme. The addressing scheme may conform to

the methodology described in RFC 3177 "IAB/IESG Recommendations on IPv6 Address Allocations to Sites".

Allocating global addresses manually has many drawbacks. In particular, it takes a team of specialist technicians a long time. Nor does it provide for easy reconfiguration of the topology of the network or the addition of new router into a pre-existing network. Most importantly, it may lead to human errors, however competent the technicians may be. The larger the network, the greater the number of errors, making them difficult to detect and correct.

A need has become apparent for automating the process of configuring communications networks and in particular for allocating interface global addresses.

A first step toward such automation is disclosed in the documents "Automatic Prefix Delegation Protocol for Internet Protocol Version 6 (IPv6)" by B. Haberman and J. Martin (draft-haberman-ipngwg-auto-prefix-02.txt), published February 2002, and "Hierarchical Prefix Delegation Protocol for Internet Protocol Version 6 (IPv6) by Byung-Yeob Kim, Kyeong-Jin Lee, Jung-Soo Park and Hyoung-Jun Kim (draft-bykim-ipv6-hpd-00.txt) published October 2003. The above two documents are IETF (Internet Engineering Task Force) drafts available with the file names given in brackets above on the IETF web site.

The above documents disclose automatically allocating an address prefix to a router, which address prefix is based on an address prefix supplied by another router. That other router is called the "address delegator" and this mechanism is called "address delegation".

A mechanism of the above kind is insufficient, however, since it merely allocates an address prefix of the router. There is no reference in either of the above two documents to allocating global addresses to the interfaces of the routers. Unfortunately, such

allocation is necessary in order to achieve total automation of network configuration, and allocating interface global addresses encounters problems that the above documents pass over in silence.

The technical problem addressed by the invention is therefore that of allocating a unique global address to each of the interfaces of each of the routers of an IPv6 communications network.

To this end, the invention provides a communications router for an Internet communications network (in particular an Ipv6 communications network), the router comprising a set of interfaces each connected to one or more other communications routers. The router (or network equipment) comprises means for receiving an address prefix from a first other communications router over a first interface. The communications router of the invention is characterized in that it also comprises allocation means for allocating each of said interfaces a global address determined in particular from said address prefix.

In one embodiment of the invention, the allocation means determine the global address of one of the interfaces by concatenating an interface identifier with a network number containing the address prefix and forming an addressing sub-space of the addressing space formed by the address prefix.

In a preferred embodiment of the invention, the allocation means allocate the first interface the same network number as the first communications router allocates the interface connected to the first interface.

In one embodiment of the invention only one network number is allocated for each connection.

A router of the invention is therefore capable of configuring its interfaces with a unique global address, i.e. a global address that is unique for the whole of the communications network. The configuration process is automatic for each router and the network may therefore

be configured recursively and entirely automatically, starting from an initial prefix, which may be allocated automatically or otherwise.

The invention and other advantages of the invention become more clearly apparent in the course of the following description with reference to the accompanying drawings, in which:

Figure 1, already commented on, is a diagram of a communications network comprising four routers;

Figure 2 shows the format of an interface global address of the invention; and

Figure 3 is a diagram showing a communications network and the use in that network of the method described herein.

An interface global address of the invention comprises four parts, as shown in Figure 2. In an embodiment of the invention using the IPv6 protocol, the total size of the global address is 128 bits.

The part U on the extreme right is formed from a universal identifier, as is known in the art and explained above. In the IPv6 implementation, the size of this field U is 64 bits.

The part P on the extreme left is a prefix that is supplied by another router of the communications network. As explained below, the size of the prefix varies and depends on the position of the router in the topology of the communications network during the address delegation process.

The part N is a number used to split the unique resource P into a plurality of smaller resources.

The next part Z may consist of only "0" bits, for example. Its size is imposed by the size of the other fields and may in the limit be reduced to zero bits.

As a general rule, after it is started up, if it has no global prefix, a router sends a request over all of its interfaces in order to obtain an address prefix P. It has to wait until it receives this prefix because

otherwise it cannot determine the global addresses of its interfaces.

As soon as it has received a prefix, it may resume the process of allocating interface global addresses. Arbitration may be required if it receives a plurality of prefixes from a plurality of other routers connected via its interfaces. For example, the choice may be made to take the first prefix received and to consider the router from which it comes as the "delegator" for the remainder of the process.

Once the prefix has been obtained, the router can determine a global address for all its interfaces; in a preferred embodiment of the invention this excludes the global address of the "delegator", i.e. of the interface connected to the router that supplied it with the prefix. As explained later, this interface may be judiciously configured with the value of the received prefix, which guarantees good aggregation of routes in routers on the upstream side of the delegator.

The global addresses are determined in accordance with the following scheme, using the received prefix to construct the field P and using the universal identifier to construct the field U.

In one embodiment of the invention, the router does not determine the global address of the interface through which the prefix was sent to it.

The router determines a set of sub-prefixes SP defining an address space that is smaller than that defined by the prefix P. These sub-prefixes SP therefore contain the prefix P and concatenate with it on its right-hand side a certain number of bits (this is the part N).

These sub-prefixes SP and the associated addressing spaces may be used by the router for its own purposes or they may be delegated to other communications routers.

The sub-prefixes SP may be determined from the prefixes P in various ways.

The table below lists the sub-prefixes SP that may be formed from a prefix P. In this example, the choice has been made to concatenate three additional bits to the left-hand side of the prefix to form the sub-prefixes. There are therefore eight sub-prefixes, because $2^3 = 8$. The first column gives the sub-prefixes in hexadecimal and the second column gives them in binary.

P-0000	P-0000 0000 0000 0000
P-2000	P-0010 0000 0000 0000
P-4000	P-0100 0000 0000 0000
P-6000	P-0110 0000 0000 0000
P-8000	P-1000 0000 0000 0000
P-A000	P-1010 0000 0000 0000
P-C000	P-1100 0000 0000 0000
P-E000	P-1110 0000 0000 0000

In this example, the choice has been made that each sub-prefix should be the same size (equal to the size of the prefix P + 3 bits), forming as many addressing sub-spaces of equal size. Other implementations are possible, of course.

There is also more than one way to distribute the sub-prefixes. For example, there is the delegation process described above and specified in more detail in the documents cited above. In this case the router in question functions as a prefix delegator in relation to other routers connected to its interfaces.

In one particular embodiment of the invention, the use of the addressing capacity of the global addresses may be optimized by choosing the length to add to the prefix received as a function of the number of routers connected.

Accordingly, if a router has "n" routers connected to it, then $\log_2(n)$ bits are required to represent the router number, where $\log_2$ denotes the logarithm to base 2. The addressing space linked to this number "n" may be used to number the neighboring routers, i.e. those connected directly to its interfaces. For example, if

the router has three neighboring routers, it may number them on 2 bits ($\log_2(3) = 2$), and they could take the identification numbers 1, 2 and 3 (respectively "01", "10", and "11" in binary).

Given that a neighboring router is accessible via an interface of the router in question, one embodiment of the invention proposes to number a given interface as a function of the prefix allocated to it from the neighboring router. In Figure 1, for example, the router R_A allocates the same number to its interface I_{A2} as the prefix that it has allocated to the router R_D.

With regard to the interface I_{A1}, arbitration may be required between the numbers allocated to the routers R_B and R_C, as they are connected to the same interface I_{A1}.

Thus the above process determines at least one global address that is the same for the interfaces of routers that are connected together, thereby conforming to the IPv6 protocol.

Determining at most one address is highly beneficial as it economizes on addressing resources, because in this case the same network number is used for three interfaces, which saves two network numbers that may therefore be used for other purposes.

Accordingly, in Figure 1, the network numbers (i.e. the 64 more significant bits of the global addresses) of the interfaces I_{A1}, I_B, and I_C are the same.

Figure 3 shows the use of the method of the invention in a small network.

The router R_1 has acquired an initial prefix with the value 2001:db8:1:0000::0/48. The meaning of this format is explained in the documents on IPv6 protocol address formats cited above. However, it is important to note here that the "48" indicates the length in bits of this prefix, which is 64 bits maximum. Below, with respect to interface global addresses, there is no reference to the universal identifier part U.

This initial prefix is used by the router R_1 to

determine the sub-prefixes that it is going to send to its neighboring routers in its capacity of delegator. In this example, three additional bits are allocated to divide the addressing space into $2^3 = 8$ smaller spaces (the size of the field N is 3 bits). Thus a first sub-space, for example "001", is allocated to the router R_2 and a second sub-space, for example "010", is allocated to the router R_5. The sub-prefixes sent to the routers R_2 and R_5 are 2001:db8:1:2000::0/51 and 20001:db8:1:4000::0/51, respectively.

The router R_1 allocates network numbers to its interfaces as a function of the sub-prefixes allocated to the routers connected to those interfaces.

The network number 2001:db8:2000::0/64 is therefore obtained for the interface I_{1a}. As indicated above, this 64-bit network number may be completed by adding the universal identifier to obtain the 128 bits of the IPv6 interface global address, for example.

Similarly, the 64-bit network number 2001:db8:1:4000::0/64 is obtained for the interface I_{1b}.

The routers R_2 and R_5 thereafter proceed in the same manner, for example as soon as they are in possession of the prefix sent by the delegator R_1. The process may therefore be executed partly in parallel on the various routers.

The router R_2 allocates the interface I_{2a} the same network number that the router R_1 allocates the interface I_{1a}.

In the same manner as before, the router R_2 allocates three additional bits to divide into eight parts its addressing space determined by the prefix supplied by the delegator R_1.

The router R_5 is allocated a sub-prefix by the router R_1. It is at the same level in the delegation tree as the router R_2, so to speak. Consequently, the router R_2 does not take it into consideration when allocating a sub-prefix.

The router R_2 therefore allocates the router R_4 a first sub-space (for example "001" (N.B. these digits are the values of the bits of the field N)) and allocates the router R_5 a second sub-space (for example "010").

The sub-prefixes that are sent to them are therefore 2001:db8:1:2400::0/64 and 20001:db8:1:2800::0/64, respectively, as "2400" is "001-001-00…" in binary and "2800" is "001-010-00…" in binary. Note that the first group on the left-hand side ("001") consists of the values of the three bits allocated to the router R_2 by the router R_1.

The network number of the interfaces I_{2c} and I_{5c} cannot be allocated in the same way and is therefore negotiated between the two routers R_2 and R_5. In the Figure 3 example, the result of this negotiation is that this number must be determined by the router R_5. Consequently, the network number of the interfaces I_{2c} and I_{5c} contains the sub-prefix of the router R_5, on the same terms as the network number of the interface R_{5a}, for which the situation is exactly the same.

In this example, the router R_5 therefore allocates three additional bits to divide its addressing space into eight portions. It therefore allocates a first sub-space (for example "001") for the interface I_{5a} and a second sub-space (for example "010") for the interface I_{5c}. The network numbers of these two interfaces are then 2001:db8:1:4400::0/64 for the interface I_{5a} and 2001:db8:1:4800::0/64 for the interface I_{5c}.

The situation is different for the interfaces I_{2b}, I_{3b} and I_4, as there is a delegation relationship between the router R_2, on the one hand, and the routers R_3 and R_4, on the other hand. Once again, negotiation may be required to determine if it is the sub-prefix allocated to the router R_3 or that allocated to the router R_4 that must be used. This negotiation means that a single network prefix can be used for all the interfaces, although a different embodiment could choose different

network addresses, either in accordance with some other negotiation mechanism or without negotiation.

In the Figure 3 example, the sub-prefix allocated to the router R_4 is chosen. The network number of these three interfaces I_{2b}, I_{3b}, and I_4 is therefore 2001:db8:1:2400::0/64.

An additional advantage of the invention is that, because the process is tree-structured, each router of the network allocates global addresses to its interfaces that are formed on the basis of a prefix supplied by the delegator. In other words, from the point of view of the delegator, all the network numbers and therefore all the global addresses of the interfaces of the routers to which it has supplied a prefix are "aggregatable" addresses of that prefix (the expression "aggregatable addresses" means addresses constructed from the same prefix). Also, these "aggregatable addresses" may be stored in the form of a single entry in the routing table of the delegator router.

This saves memory space in the router and saves time looking up the correct entry in the routing table if the router has to route data packets.